

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
3404 E. Harmony Road
Mail Stop 35
Fort Collins, Colorado 80528

PATENT APPLICATION

ATTORNEY DOCKET NO. 200313629-1

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Simon C. Steely Jr. et al.

Confirmation No.: 9866

Application No.: 10/760,659

Examiner: Mardochee Chery

Filing Date: January 20, 2004

Group Art Unit: 2186

Title: System and Method for Non-Migratory Requests in a Cache Coherency Protocol

Mail Stop Appeal Brief - Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF REPLY BRIEF

Transmitted herewith is the Reply Brief with respect to the Examiner's Answer mailed on 04/29/2010 .

This Reply Brief is being filed pursuant to 37 CFR 1.193(b) within two months of the date of the Examiner's Answer.

(Note: Extensions of time are not allowed under 37 CFR 1.136(a))

(Note: Failure to file a Reply Brief will result in dismissal of the Appeal as to the claims made subject to an expressly stated new ground rejection.)

No fee is required for filing of this Reply Brief.

If any fees are required please charge Deposit Account 08-2025.

Respectfully submitted,
Simon C. Steely Jr. et al.

By: /Steven L. Nichols/
Steven L. Nichols
Attorney/Agent for Applicant(s)

Reg No. : 40,326

Date : June 7, 2010

Telephone : 801-237-0251

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In the Patent Application of

Simon C. Steely, Jr. et al.

Application No. 10/760,659

Filed: January 20, 2004

For: System and Method for Non-Migratory
Requests in a Cache Coherency Protocol

Group Art Unit: 2186

Examiner: Chery, Mardochee

Confirmation No.: 9866

REPLY BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is a Reply Brief under Rule 41.41 (37 C.F.R) in response to the Examiner's Answer of April 29, 2010 (the "Examiner's Answer" or the "Answer"). In Section 10, the Answer contains a response to some of the arguments made in Appellant's brief. Appellant now responds to the Examiner's Answer as follows.

Status of Claims

Claims 2-7 and 13-17 have been identified as reciting allowable subject matter, but stand objected to as being dependent upon a rejected base claim.

Claims 1, 8-12, and 18-32 are pending in the application and stand finally rejected.

Accordingly, Appellant appeals from the final rejection of claims 1, 8-12, and 18-32, which claims are presented in the Appendix.

Grounds of Rejection to be Reviewed on Appeal

The final Office Action raised the following grounds of rejection.

(1) Claims 1, 8-10, 12, 18-19, and 24-32 were rejected under 35 U.S.C. § 103(a) as being obvious over U.S. Patent Publication No. 2002/0129211 by Arimilli et al. (“Arimilli”) in view of U.S. Patent No. 6,931,496 to Chen et al. (“Chen”).

(2) Claims 11 and 20-23 were rejected under 35 U.S.C. § 103(a) as being obvious over Arimilli in view of Chen and further in view of U.S. Patent No. 6,484,240 to Cypher et al. (“Cypher”).

Accordingly, Appellant has requested review of each of these grounds of rejection in the present appeal.

VII. Argument

(1) Claims 1, 8-10, 12, 18-19, and 24-32 are patentable over Arimilli and Chen:

Claims 1, 12, 24, 28, and 32:

In the present application, independent claim 1 recites:

A system comprising:

a first node including data having an associated state, the associated state of the data at the first node being a modified state; and

a second node operative to provide a non-migratory source broadcast request for the data, the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating memory, the second node being operative to receive the data from the first node and assign a shared state to an associated state of the data at the second node.

(Emphasis added).

Independent claim 12 recites:

A multi-processor network comprising:

memory for storing data;

a first processor node having a first processor node cache line including the data, the first processor node cache line having an associated state, the associated state of the first processor node cache line being a modified state; and

a second processor node operative to provide a non-migratory source broadcast read request for the data, the second processor node having a second processor node cache line with an associated state;

the first processor node being programmed to respond to the non-migratory source broadcast read request of the second processor node by providing a shared data response to the second processor node and transitioning the associated state of the first processor node cache line from the modified state to an owner state without updating the memory with the data, the data being stored in the second processor node cache line, the associated state of the second processor node cache line being assigned a shared state.

(Emphasis added).

Independent claim 24 recites:

A system comprising:
means for broadcasting from a first node a non-migratory read (XREADN) request for data;
means for providing the *data from a second node to the first node in response to the XREADN request, a modified state being associated with the data at the second node, a shared state being associated with the data at the first node in response to the first node receiving the data from the second node*; and
means for *transitioning the modified state associated with the data at the second node to an owner state without updating memory of the system.*
(Emphasis added).

Independent claim 28 recites:

A method comprising:
broadcasting a non-migratory request for data from a first node to other nodes of an associated system;
providing a shared copy of the data from a second node to the first node in response to the non-migratory request;
transitioning a state associated with the data at the second node from a modified state to an owner state in response to the non-migratory request; and
transitioning a state associated with the data at the first node to a shared state in response to receiving the shared copy of the data from the second node.
(Emphasis added).

Independent claim 32 recites:

A computer system comprising a cache coherency protocol that is operative to permit migration of data to a cache associated with a source processor from a cache associated with a target processor *when a migratory request is issued from the source processor, the protocol being further operative to prevent migration of the data to the cache associated with the source processor from the cache associated with the target processor when a non-migratory request is issued from the source processor.*
(Emphasis added).

Attacking References Individually

Appellant has demonstrated in the Brief dated February 10, 2010 that the Office has failed to meet its burden to establish a *prima facie* case of obviousness against the subject

matter of each of the above independent claims because the Arimilli and Chen references do not teach or suggest, individually or jointly, all of the subject matter recited in the independent claims. The Answer is quick to assert that Appellant's Brief merely attacks references individually, and that "one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references." (Answer, p. 12) (citing to *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Appellant wishes to respectfully point out that where a rejection relies on an assertion of obviousness based on the combination of prior art elements, the M.P.E.P. requires Office personnel to resolve the . . . factual inquiries" of *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17-18 (1966). M.P.E.P. § 2143. Under the analysis required by *Graham*, to support a rejection under § 103, the scope and content of the prior art must first be determined, followed by an assessment of the differences between the prior art and the claim at issue in view of the ordinary skill in the art. *Graham*, 383 U.S. at 17-18. The scope and content of the prior art as a whole may only be determined by a methodical examination of *each* prior art reference to determine what each reference teaches and fails to teach. If such an examination reveals that *each* cited prior art reference fails to teach or suggest a particular claimed element, it follows that the scope and content of the cited prior art does not include that element.

Moreover, the Answer relies on a misreading of the *Keller* and *Merck* cases to support the Examiner's contention that any examination of the prior art references individually is irrelevant to a determination of whether a claim is obvious over the combination of those prior art references. *Keller* and *Merck* do not support such a proposition.

In the *Keller* case, a claim was rejected as obvious over an issued patent in view of a published journal article, and the Court of Customs and Patent Appeals held that the applicant did not sufficiently rebut a *prima facie* case of obviousness by introducing an affidavit that attacked the teachings of only the published journal article. *Keller*, 642 F.2d at 425-26. In the *Merck* case, the Federal Circuit found that an attack on a single prior art reference was insufficient to overcome an established *prima facie* case of obviousness based on a combination of many prior art references. *Merck*, 800 F.2d at 1094, 1099. Thus, the *Keller* and *Merck* cases simply stand for the fact that where a *prima facie* case of obviousness has been established against a claim, the *prima facie* obviousness of that claim cannot be rebutted by attacking *only one* of the prior art references used to establish the *prima facie* case of obviousness.

The present claims at appeal are distinguishable from the *Keller* and *Merck* cases at least because a) no *prima facie* case of obviousness has been established by the Office against Appellant's rejected claims in the present case, and b) Appellant has demonstrated in the present case the failure of *both* cited prior art references to teach various claimed elements (as opposed to attacking only *one* of the cited prior art references).

As such, Appellant's demonstration that each cited prior art reference does not teach or suggest various elements recited in the independent claims is simply an evaluation of the "scope and content of the prior art" according to the first prong of the *Graham* inquiry. Where the scope and content of the prior art does not include elements recited in a claim, the burden is on the Office to demonstrate that those elements would be obvious to one having ordinary skill in the art at the time of the invention in order to establish a *prima facie* case of obviousness against the claim. M.P.E.P. § 2143.

Functional Limitations in System Claims

In various portions of the Answer, the Examiner asserts that the functional limitations recited in a system claim do not have patentable weight because “claims directed to a system must be distinguished from the prior art in terms of structure rather than function.” (Answer, pp. 13-14) (citing to *In re Schreiber*, 128 F.3d 1473, 1477-78, 44 USPQ2d 1429, 1431-32 (Fed. Cir. 1997)). However, this position is a misreading of the applicable case law. *Schreiber* never states that an apparatus must be distinguished from the prior art in terms of structure rather than function to be patentable. Rather, *Schreiber* holds that functional limitations of a claim do not give the claim patentable weight **only if** the limitations are “found to be inherent in . . . the prior art reference,” and that where functional language is the only distinguishing feature between a claim and the cited prior art, the burden is on the Applicant to “show that the prior art structure did not inherently possess the functionally defined limitations of his claimed apparatus.” *Schreiber*, 128 F.3d at 1478.

The other cases cited to by the Examiner are very clear in their support of the position that functional subject matter in a system claim **does** carry patentable weight. *In re Swinehart* states that

where the Patent Office has reason to believe that a functional limitation asserted to be critical for establishing novelty in the claimed subject matter may, in fact, be an inherent characteristic of the prior art, it possesses the authority to require the applicant to prove that the subject matter shown to be in the prior art does not possess the characteristic relied on.

439 F.2d 210, 212, 169 USPQ 226, 228 (CCPA 1971). As such, *Swinehart* expressly indicates that functional limitations in system claims carry patentable weight and may be “critical for establishing novelty.” *Id.*

The court in *In re Danly* held that where apparatus claims recite functional limitations with “the obvious intention of limiting them to actual performance of the stated functions, as

distinguished from mere possibility of such performance,” the functional limitations are sufficient to distinguish the claims over the prior art in order to allow the claims. 263 F.2d 844, 847, 120 USPQ 528, 531 (CCPA 1959).

In *Hewlett-Packard Co. v. Bausch & Lomb Inc.* no functional language in a system or apparatus claim was even at issue. 909 F.2d 1464, 15 USPQ2d 1525, 1528 (Fed. Cir. 1990). Rather, the Federal Circuit held that an applicant has no obligation to demonstrate a difference in functionality over the prior art for a system or apparatus claim that is structurally different from the prior art because “apparatus claims cover what a device *is*, not what a device *does*.” *Id.* at 1468 (emphasis original). Continuing, the Court held that “[a]n invention need not operate differently than the prior art to be patentable, but need only be different.” *Id.* However, in making this statement the Court did not say that functional limitations have no patentable weight, only that functional distinctions over the prior art are *unnecessary* to patentability where structural differences exist.

Appellant’s initial Appeal Brief and the present Reply Brief provide various reasons why the all of functionality described in Appellant’s independent claims is not expressly taught or suggested or inherently present in the cited Arimilli and Chen references. This failure by the prior art references to teach or suggest all of the claimed subject matter cannot be ignored in making an inquiry as to the novelty of the claimed invention. M.P.E.P. § 2173.05(g) (“A functional limitation **must** be evaluated and considered, **just like any other limitation of the claim**, for what it fairly conveys to a person of ordinary skill in the pertinent art in the context in which it is used”) (emphasis added); *see also* M.P.E.P. § 2114 (“features of ***an apparatus*** may be recited either structurally ***or functionally***”) (citing to *Swinehart*, 439 F.2d at 212).

Thus, the Examiner's comments regarding the functional recitations in Appellant's claims are misplaced and improper. Furthermore, for at least the reason that Arimilli and Chen do not teach or suggest all of the recited functionality of the independent claims, the rejection of Appellant's independent claims should not be sustained.

"Second Node"

Appellant has taken the position in the Appeal Brief that Arimilli and Chen fail to teach or suggest a second node that is "operative to receive data from the first node" and "operative to assign a shared state to an associated state of the data at the second node." (Brief, pp. 13-14) (citing to claim 1). The Answer discounts this position as "clearly erroneous" and an "assumption" in light of paragraph 0029 of Arimilli. (Answer, pp. 12-13).

The portion of this paragraph cited by the Examiner reads as follows:

Master 26 of processor complex 10a stores to the target cache line and completes associated clean-up operations, if any, (also discussed below) at a later time t2. The interval 160 between time t0 and t2 represents the time window in which a conflict between modifying transaction 150a and one or more other requests to modify the target cache line may arise. As defined herein, a conflict arises if, during interval 160, the master 26 of a second agent 10 (and possibly one or more additional agents 10) develops, or has previously developed and manifests, at any time during interval 160 an intention to modify the target cache line. This intention, which arises as a result of a store request by the processor 16 in the second agent 10 that targets a cache line marked as shared in the second agent's cache directory 22

According to the Examiner, the "processor complex 10a" reads on the first node recited in claim 1 and the "agent 10" reads on the second node recited in claim 1. (Answer, pp. 12-13, 14-15). Appellant respectfully disagrees. In the above scenario of Arimilli, the only data received by the "processor complex 10a" from the agent 10 are conflicting requests to modify the target cache line. Thus, for the processor complex 10a to be "operative to assign a shared state to an associated state of *the* data" received from the agent 10, Arimilli

must teach or suggest that the processor complex 10a assigns a shared state to a request to modify the target cache line received from the agent 10. Arimilli does not teach or suggest this subject matter anywhere. By contrast, Arimilli merely teaches that data intended to be modified by more than one agent is marked as shared in response to a received request to modify the data.

The Answer further cites to Chen's teaching of first and second nodes where "[t]he first node includes an external cache for storing a data from a local memory of the second node and at least two processors optionally accessing the data from the external cache." (Answer, p. 13) (citing to Chen, abstract). However this portion of Chen does not teach or suggest that either of the nodes in Chen is assigns a shared state to an associated state of the data at the other node.

In light of these facts, Appellant respectfully submits that the Answer and final Office Action do not make a *prima facie* case of obviousness against the "second node" recited in Appellant's independent claims. Consequently, the rejection of Appellant's independent claims should not be sustained for at least this reason.

"First Node"

In response to Appellant's position that Arimilli and Chen both fail to teach or suggest a "first node [that] is operative to provide the data to the second node and transition the associated state of the data at the first node from the modified state to an owner state in response to a non-migratory source broadcast request provided by the second node," the Answer cites to various portions of Appellant's Specification that teach the claimed first node. (Answer, pp. 16-18) (citing to Appellant's Specification, paras. 0037-39). Appellant wishes to respectfully remind the Examiner that Appellant's Specification is not prior art

against Appellant's claims. Furthermore, "it is important not to import into a claim limitations that are not part of the claim. For example, a particular embodiment appearing in the written description may not be read into a claim when the claim language is broader than the embodiment." *Superguide Corp. v. DirecTV Enterprises, Inc.*, 358 F.3d 870, 875, 69 USPQ2d 1865, 1868 (Fed. Cir. 2004).

Regarding the "first node" recited in claim 1, the Answer further cites to Chen's teaching of a local node that stores data "exclusively owned by a certain processor of the local node" such that when "a command is issued to read an exclusive copy of the specific data which is stored in the local node," the local node "chang[es] the state of the local memory line of the remote node where the specific data is stored, from GONE into HOME, or chang[es] the state of the local memory line of the remote node, from SHARED to HOME." (Answer, p. 19) (citing to Chen, col. 5 lines 5-10 and 36-42, col. 5 line 60 to col. 6 line 6). Nevertheless, this portion of Chen refers only the marking of a "local memory line," and not the marking of the data itself. Furthermore, this marking does not occur "in response to a non-migratory source broadcast request," as the remote node takes ownership of the data requested from the local node. (Chen, col. 5 lines 5-10 and 36-42, col. 5 line 60 to col. 6 line 6).

The Answer further cites to teachings in Chen of "changing the state of the cache line from the CLEAN state to either of DIRTY-SHARED and DIRTY-ONLY state if the data has been modified into a modified data and allowing a second certain one of the at least two processors to directly request the modified data via a bus inside the first node when the cache line is in the DIRTY_SHARED state." (Answer, pp. 19-22) (citing to Chen, col. 3 lines 55-65, col. 6 line 48 to col. 7 line 2). Nevertheless, this portion of Chen does not teach "transition[ing] the associated state of the data at the first node from the modified state to an

owner state in response to a non-migratory source broadcast request provided by the second node.” (Claim 1). Rather, this portion of Chen deals with changing the state of a *cache line* and not the individual data stored within that cache line. Furthermore, none of the CLEAN, DIRTY-SHARED, and DIRTY-ONLY states described by Chen is analogous to “an ownership state” recited in Appellant’s independent claims.

In response to Appellant’s position that Chen fails to teach or suggest a first node that “transition[s] the associated state of the data at the first node from the modified state to an owner state *without updating memory*,” as recited in claim 1, the Answer cites to Chen’s Figs. 2-5. (Answer, p. 22). However, it will be readily apparent to anyone having ordinary skill in the art that Chen does not teach or suggest in Figs. 2-5 or anywhere else that the state of data at a first node is modified “without updating memory.” (Claim 1).

Accordingly, the Office has not demonstrated that the “first node” recited in claim 1 is taught or suggested by the Arimilli and Chen references. As such, the Office has not met its burden to establish a *prima facie* case of obviousness against Appellant’s independent claims. Therefore, the rejection of the independent claims should not be sustained for at least this additional reason.

Motivation to Combine Arimilli and Chen

Appellant’s initial Brief points out that one having ordinary skill in the art would have found no motivation to combine the Arimilli and Chen references at the time of the invention because Chen is directed to a distributed shared memory (DSM) system and Arimilli is directed to a central memory system, and the transactions required to maintain coherency in a DSM system would adversely affect the principal operation of a central memory system. (Brief, p. 18). The Answer dismisses this position as “mere allegations for there is nothing

either in Arimilli or in Chen supporting appellant's position and combining Arimilli and Chen would provide no adverse result." (Answer, p. 24). Appellant respectfully disagrees.

The Examiner's position that in order for there to be no motivation to combine two references, evidence of the lack of motivation must exist within the references themselves is circular and logically untenable. Those having ordinary skill in the art will readily recognize that a DSM system, in which shared data is distributed across multiple local node caches (Chen, Fig. 1 and col. 1 line 63 to col. 2 line 25), is functionally and operationally different from a central memory system, in which shared data is stored in a system memory that is independently addressed on the network (Arimilli, Fig. 1 and paras. 0004-10).

Arimilli is principally directed to the arbitration of conflicting requests to modify data in a shared system memory (Arimilli, para. 0011), whereas Chen is directed to "provid[ing] a data maintenance method of a DSM system to achieve the system bus transaction purpose with reduced transaction requests" that "solve[s] the access dead lock problem." (Chen, col. 3 lines 6-12). Thus, where Arimilli solves the problem of selecting between competing nodes that wish to modify data in a stored central memory, Chen solves the problem of reducing bus congestion among nodes of a DSM system. These problems are substantially unrelated, and therefore a person having ordinary skill in the art would not have looked to Chen as a source of viable modifications to Arimilli.

"A patent [or patent application claim] composed of several elements is not proved obvious merely by demonstrating that each element was, independently, known in the prior art." *KSR Int'l Co. v. Teleflex, Inc.*, 550 U.S. 398, 418 (2007). According to the Supreme Court, "it can be important to identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does. This is so because inventions in most, if not all, instances rely upon building

blocks long since uncovered, and claimed discoveries almost of necessity will be combinations of what, in some sense, is already known.” *Id.* at 418-19.

“To determine whether there was an apparent reason to combine the known elements in the way a patent claims, it will often be necessary to look to interrelated teachings of multiple patents; to the effects of demands known to the design community or present in the marketplace; and to the background knowledge possessed by a person having ordinary skill in the art. To facilitate review, this analysis should be made explicit.” *Id.* at 418.

In light of the standard imposed by *KSR*, Appellant respectfully submits that the Office has failed to provide an explicit analysis establishing how a person having ordinary skill in the art would have been motivated to combine the Arimilli and Chen references. Accordingly, the combination of Arimilli and Chen to reject Appellant’s independent claims is improper. Consequently, the rejection of Appellant’s independent claims should not be sustained for at least this additional reason.

Claims 8, 10, 19, 25-27, and 29-31:

Appellant maintains that no *prima facie* case of obviousness has been established against dependent claims 8, 10, 19, 25-27, and 29-31 for at least the same reasons given in the Appeal Brief and above in favor of the patentability of the independent claims from which they respectively depend, and for the reasons given in the Appeal Brief in favor of the patentability of the additional subject matter recited in each of these dependent claims. Additionally, Appellant’s response to the Answer’s arguments with respect to attacking references individually and functional recitations against the independent claims is equally applicable to the same arguments made by the Answer against claims 8, 10, 19, 25-27, and 29-31.

(2) Claims 11 and 20-23 are patentable over Arimilli, Chen, and Cypher:

Appellant maintains that no *prima facie* case of obviousness has been established against dependent claims 11 and 20-23 for at least the same reasons given in the Appeal Brief and above in favor of the patentability of the independent claims from which they respectively depend, and for the reasons given in the Appeal Brief in favor of the patentability of the additional subject matter recited in each of these dependent claims. Additionally, Appellant's response to the Answer's arguments with respect to attacking references individually and functional recitations against the independent claims is equally applicable to the same arguments made by the Answer against claims 11 and 20-23.

In view of the foregoing, it is again submitted that the final rejection of the pending claims is improper and should not be sustained. Therefore, a reversal of the Rejection of December 10, 2009 is respectfully requested.

Respectfully submitted,

DATE: June 7, 2010

/Steven L. Nichols/
Steven L. Nichols
Registration No. 40,326

Steven L. Nichols, Esq.
Vancott, Bagley, Cornwall & McCarthy
36 South State Street
Suite 1900
Salt Lake City, Utah 84111

(801) 237-0251 (phone)
(801) 237-0871 (fax)